

Bounding Optimality Gaps for Non-Convex Optimization Problems: Applications to Nonlinear Safety-Critical Systems

Prithvi Akella and Aaron D. Ames¹

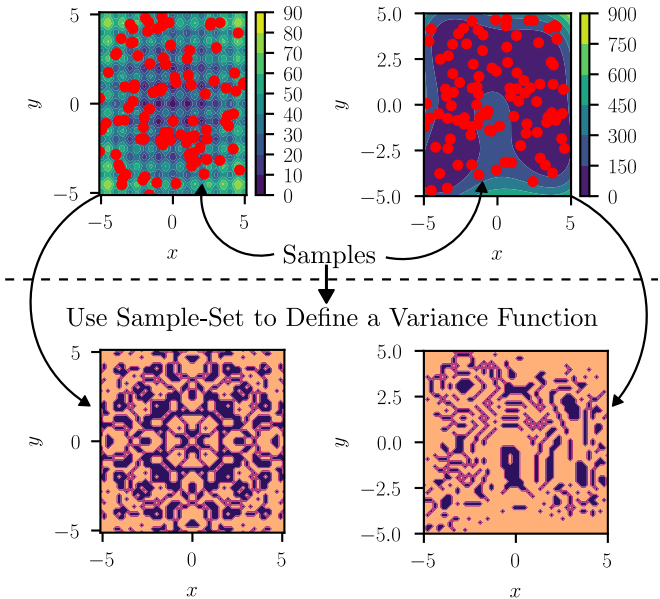
Abstract—Efficient methods to provide sub-optimal solutions to non-convex optimization problems with knowledge of the solution’s sub-optimality would facilitate the widespread application of nonlinear optimal control algorithms. To that end, leveraging recent work in risk-aware verification, we provide two algorithms to (1) probabilistically bound the optimality gaps of solutions reported by novel percentile optimization techniques, and (2) probabilistically bound the maximum optimality gap reported by percentile approaches for repetitive applications, *e.g.* Model Predictive Control (MPC). Notably, our results work for a large class of optimization problems. We showcase the efficacy and repeatability of our results on a few, benchmark non-convex optimization problems and the utility of our results for controls in a Nonlinear MPC setting.

I. INTRODUCTION

Optimal controllers have emerged in recent years as the preeminent choice for controller synthesis as they offer a natural way of expressing (perhaps) disparate control objectives [1]–[3], *e.g.* as in Model Predictive Control (MPC) [4]–[6], control-barrier-function based quadratic programs [7]–[9], and optimal path planning [10]–[12], among others. Specifying to the first two examples, optimal controllers in this vein leverage existing system models to inform controller choice at the current or successive time-steps, *e.g.* the predictive aspect of MPC references the utilization of a system model to predict how a given input will affect system evolution and uses that prediction to determine optimal input choice for the next few time-steps. Notably, one-step control-barrier-function-based quadratic programs ensure convexity for nonlinear system models, though predicting over these models remains a non-convex problem [13]–[16].

The utility of these controllers inspired efforts to develop and exploit rapid solution techniques for these non-convex optimization problems, though some improvement avenues still exist. For example, typical approaches in the literature include Model Predictive Path Integral Control (MPPI) [17], [18], Collocation methods [19], [20], Bayesian Optimization [21], [22], Learning-based methods [23], [24], and more recently, percentile-based methods [25], [26]. In specific for the latter three learning-based methods, both bayesian optimization and other learning-based methods guarantee eventual convergence to an optimal solution but do not provide sample complexity bounds. Percentile methods offer robust sample complexity bounds and have seen considerable success at efficiently providing sub-optimal solutions to

Approximate Optimal Solutions via Sampling Methods



Upper-Bound Optimality Gap via Sampling Variances

Fig. 1. An overview of our proposed method. Sampling methods employed in risk-aware verification provide decent results to non-convex optimization problems, though their optimality gaps are not well-understood. By defining a variance function around the previously generated sample set, we can implicitly identify decisions (orange on the bottom) whose variances exceed the optimality gap of our reported sampling solution and identify such an optimality gap via a second application of the same sampling methods.

constrained nonlinear optimal controllers. However, unlike the first two methods, they fail to provide even probabilistic bounds on the optimality gap between the reported percentile solution and the true optimal value.

Our Contribution: We aim to provide probabilistic guarantees on the optimality gap of percentile solutions — the deviance in cost evaluation of the percentile solution and the minimum function value — for their respective optimization problems. In doing so, we hope to develop the first in a line of algorithms that provide robust sample complexity bounds for efficient identification of sub-optimal solutions with knowledge of the solution’s sub-optimality. To that end, our contributions are twofold:

- Subject to a fairness assumption, we provide sample-complexity bounds for the identification of the optimality gap for percentile solutions to a large class of optimization problems.
- For optimization problems to be solved repeatedly, *e.g.* MPC, we provide a method to determine the probability of percentile solutions exhibiting optimality gaps lower

*This work was supported by the AFOSR Test and Evaluation Program, grant FA9550-19-1-0302

¹All authors are with the California Institute of Technology {pakella, ames}@caltech.edu

than a determined upper bound.

We develop such a procedure to offer efficient solutions to constrained optimal control problems with known sample complexity and sub-optimality bounds. This procedure would deviate from existing approaches insofar as it would consider general (perhaps) nonlinear constraints for optimal controllers, and provide sample-complexity results on the sub-optimality of the provided solution — a discrepancy with existing methods, e.g. MPPI, Collocation, Bayesian Optimization, and other, learning-based approaches.

Structure: Section II reviews percentile optimization. Section III-A outlines our solution to bound the optimality gaps of percentile solutions. Section III-B extends these results for optimization problems to be solved repeatedly, e.g. Nonlinear MPC. Finally, Section IV provides examples to validate our theoretical results. Specifically, Section IV-A validates the probabilistic bounding results of Section III-A, and Section IV-C illustrates the utility of the results in Section III-B in a Nonlinear MPC setting.

II. A REVIEW OF PERCENTILE OPTIMIZATION

To facilitate the statement of our contributions, we first review percentile optimization in this section. Consider the following optimization problem:

$$J^* = \min_{s \in \mathbb{S}} J(s), \quad (1)$$

subject to the following assumption:

Assumption 1. The decision space \mathbb{S} is a set with bounded volume, i.e. $\int_{\mathbb{S}} 1 ds = V_{\mathbb{S}} < \infty$ or \mathbb{S} has a finite number of elements. Furthermore, the cost function J is bounded over \mathbb{S} , i.e. $\exists m, M \in \mathbb{R}$, s. t. $m \leq J(s) \leq M$, $\forall s \in \mathbb{S}$.

This assumption permits us to define two functions, \mathcal{V} , which corresponds to the volume fraction occupied by a subset A of \mathbb{S} , and F , which corresponds to the set of strictly better decisions for a provided decision $s' \in \mathbb{S}$ (we assume continuous spaces for notational ease):

$$\mathcal{V}(A) = \frac{\int_A 1 ds}{\int_{\mathbb{S}} 1 ds}, \quad (2)$$

$$F(s') = \{s \in \mathbb{S} \mid J(s) < J(s')\}.$$

Naturally then, for a given decision $s' \in \mathbb{S}$, were $\mathcal{V}(F(s')) \leq \epsilon$ for some $\epsilon \in (0, 1]$, i.e. s' is such that the volume fraction of strictly better decisions is no more than ϵ , then s' would be in the $100(1 - \epsilon)\%$ -ile with respect to minimizing J . Likewise, the associated minimum cost of such a decision $J(s')$ should also be a probabilistic lower bound on achievable costs. Both of these notions are expressed formally in the theorem below, which combines similar results from [25], [27]. To facilitate the statement of results to follow, we will first state the assumption, then the theorem.

Assumption 2. Let $I_1 = \{(s_i, J(s_i))\}_{i=1}^{N_p}$ be a set of N_p decisions and costs for decisions s_i sampled independently via $U[\mathbb{S}]$, with $\zeta_{N_p}^*$ the minimum sampled cost and $s_{N_p}^*$ the (perhaps) non-unique decision with minimum cost.

Theorem 1. Let Assumptions 1 and 2 hold. Then $\forall \epsilon \in [0, 1]$, the probability of sampling a decision whose cost is at-least $\zeta_{N_p}^*$ is at minimum $1 - \epsilon$ with confidence $1 - (1 - \epsilon)^{N_p}$, i.e.

$$\mathbb{P}_{U[\mathbb{S}]}^{N_p} \left[\mathbb{P}_{U[\mathbb{S}]} \left[J(s) \geq \zeta_{N_p}^* \right] \geq 1 - \epsilon \right] \geq 1 - (1 - \epsilon)^{N_p}. \quad (3)$$

Furthermore, $\forall \epsilon \in (0, 1]$, $s_{N_p}^*$ is in the $100(1 - \epsilon)\%$ -ile with minimum confidence $1 - (1 - \epsilon)^{N_p}$, i.e.

$$\mathbb{P}_{U[\mathbb{S}]}^{N_p} \left[\mathcal{V}(F(s_{N_p}^*)) \leq \epsilon \right] \geq 1 - (1 - \epsilon)^{N_p}.$$

III. BOUNDING THE OPTIMALITY GAP

To formally bound the optimality gap, we must first define it. To that end, consider the same general optimization problem as per (1). The optimality gap is defined as follows.

Definition 1. For general optimization problems of the form in (1), the *optimality gap* of a decision $s \in \mathbb{S}$, denoted as $G(s)$, is the deviance between the decision and optimal values, i.e. $G(s) = J(s) - J^*$.

Then our formal problem statement follows.

Problem 1. Let Assumption 2 hold, and let the optimality gap G be as per Definition 1. Identify an upper bound to $G(s_{N_p}^*)$ and the probability with which this bound holds.

A. Results for Bounded Optimization Problems

First, we note that if Assumption 2 holds, we can define the following variance function \mathbb{V} over the set I_1 :

$$\mathbb{V} : \mathbb{S} \rightarrow \mathbb{R} \text{ s. t. } \mathbb{V}(s) = \min_{s_i \in D \subseteq I_1} |J(s) - J(s_i)|. \quad (4)$$

The restriction of $s \in D \subseteq I_1$ in the definition of \mathbb{V} above is purely for practical purposes. From a theoretical standpoint, one could use the entire information set I_1 to define \mathbb{V} , but this tends to increase sample requirements as will be discussed in sections to follow (e.g. Figure 3)

Intuitively then, we aim to maximize \mathbb{V} via a percentile method to identify a variance that supersedes the optimality gap $G(s_{N_p}^*)$ of our chosen decision. To do so, we first require the following fairness assumption — that it is possible to sample variances at least as large as the optimality gap, as otherwise, it would be impossible to take a percentile approach. To formally state this assumption, we first define Ω_r to be the r -level set of \mathbb{V} :

$$\Omega_r = \{s \in \mathbb{S} \mid \mathbb{V}(s) \leq r\}. \quad (5)$$

Assumption 3. Let the variance function \mathbb{V} be as per (4), let the optimality gap G be as per Definition 1, let \mathcal{V} be as per (2), let Ω_r be as per (5), and let Assumption 2 hold. The level set $\Omega_{G(s_{N_p}^*)}$ of decisions whose variance is at most the optimality gap of $s_{N_p}^*$ does not encompass \mathbb{S} , i.e.

$$\mathcal{V}(\mathbb{S} \setminus \Omega_{G(s_{N_p}^*)}) > 0. \quad (6)$$

Second, we have the following result regarding level sets Ω .

Lemma 1. Let Ω_r be as defined in (5) and let \mathcal{V} be as per (2). The following statements are all equivalent.

- (a) $r \geq s$,
- (b) $\Omega_r \supseteq \Omega_s$,
- (c) $\mathcal{V}(\Omega_r) \geq \mathcal{V}(\Omega_s)$,
- (d) $\mathbb{P}_{\mathbb{U}[\mathbb{S}]}[\Omega_r] \geq \mathbb{P}_{\mathbb{U}[\mathbb{S}]}[\Omega_s]$.

Proof: The equivalency between (a) and (b) stems via the definition of Ω_r in (5). The equivalency between (b) and (c) stems via the definition of the volume fraction function in (2). Finally, the equivalency between (c) and (d) stems via the uniform distribution assigning probabilistic weight to subsets of $A \subseteq \mathbb{S}$ equivalent to $\mathcal{V}(A)$. ■

Third, based on our fairness assumption, we know there exists a non-zero probability of sampling decisions such that their variances are at least the optimality gap.

Lemma 2. *Let Assumption 3 hold, then there exists $p > 0$ corresponding to the probability of uniformly sampling a decision $s \in \mathbb{S} \setminus \Omega_{G(s_{N_p}^*)}$, i.e.*

$$\mathbb{P}_{\mathbb{U}[\mathbb{S}]} \left[\mathbb{S} \setminus \Omega_{G(s_{N_p}^*)} \right] = p > 0. \quad (7)$$

Proof: The result holds by definition of the uniform distribution over \mathbb{S} and equation (6). ■

Then the main result in the utilization of a percentile approach to upper bound the optimality gap stems from the prior lemmas and assumption. Formally, we aim to take a percentile approach to the following optimization problem:

$$\max_{s \in \mathbb{S}} \mathbb{V}(s), \quad (8)$$

which results in the following theorem.

Theorem 2. *Let Assumptions 1 and 3 hold, let p satisfy (7), let $I_2 = \{(s_i, \mathbb{V}(s_i))\}_{i=1}^{N_v}$ be a set of N_v decisions s_i independently sampled from $\mathbb{U}[\mathbb{S}]$ with their corresponding variances as per (4) and with $\mathbb{V}_{N_v}^*$ the maximum sampled variance. Then, $\forall \epsilon \in [0, p]$, $\mathbb{V}_{N_v}^*$ exceeds the optimality gap of the percentile solution with confidence $1 - (1 - \epsilon)^{N_v}$, i.e.*

$$\mathbb{P}_{\mathbb{U}[\mathbb{S}]}^{N_v} \left[\mathbb{V}_{N_v}^* \geq G(s_{N_p}^*) \right] \geq 1 - (1 - \epsilon)^{N_v}. \quad (9)$$

Proof: First, we know there exists a $p > 0$ satisfying (7) via Lemma 2. Second, we know that for the optimization problem (8), the decision space \mathbb{S} and objective function \mathbb{V} are bounded — this stems via Assumption 1. This permits us to take a percentile approach to solving (8). Via Theorem 1 and looking at (8) as a minimization, we have the following result substituting terms for (3) and for all $\epsilon \in [0, 1]$:

$$\mathbb{P}_{\mathbb{U}[\mathbb{S}]}^{N_v} \left[\mathbb{P}_{\mathbb{U}[\mathbb{S}]} \left[\mathbb{V}(s) \leq \mathbb{V}_{N_v}^* \right] \geq 1 - \epsilon \right] \geq 1 - (1 - \epsilon)^{N_v}.$$

The set in the interior probability corresponds to the level set of $\mathbb{V}_{N_v}^*$, i.e.

$$\mathbb{P}_{\mathbb{U}[\mathbb{S}]}^{N_v} \left[\mathbb{P}_{\mathbb{U}[\mathbb{S}]} \left[\Omega_{\mathbb{V}_{N_v}^*} \right] \geq 1 - \epsilon \right] \geq 1 - (1 - \epsilon)^{N_v}.$$

Finally, via (7), we know that the probability of sampling a decision in the level set corresponding to the optimality gap is $1 - p$. Restricting to $\epsilon \in [0, p]$ which implies that

$1 - \epsilon \geq 1 - p$ and substituting terms in the inequality above, we have the following result.

$$\mathbb{P}_{\mathbb{U}[\mathbb{S}]}^{N_v} \left[\mathbb{P}_{\mathbb{U}[\mathbb{S}]} \left[\Omega_{\mathbb{V}_{N_v}^*} \right] \geq \mathbb{P}_{\mathbb{U}[\mathbb{S}]} \left[\Omega_{G(s_{N_p}^*)} \right] \right] \geq 1 - (1 - \epsilon)^{N_v}.$$

Then, the final result holds due to Lemma 1. ■

In summary, Theorem 2 tells us that if we wish to bound the optimality gap of a percentile solution, we need to evaluate the variance of N_v uniform samples s from \mathbb{S} with respect to a subset of the information set I_1 utilized to generate the percentile solution. In practice, however, the exact probability p of sampling decisions with large enough variance will be unknown to the practitioner *a priori*. In these cases, it suffices to assume a small enough value for ϵ , i.e. 10^{-2} or smaller, is smaller than p . Examples along this vein will be provided in the following section. Notably, this result implies that we can utilize percentile methods to both identify decisions that outperform a large fraction of the decision space and also determine their optimality gap. Indeed, this result holds even for non-convex optimization problems, provided they satisfy Assumption 1.

B. Producing Solutions with Maximum Optimality Gaps

The prior section provides a method to determine the upper bound on the optimality gap of a provided solution via a secondary sampling scheme. This section provides a method to remove the secondary sampling requirement for similar optimization problems to be solved successively. In other words, consider the following general form of optimization problems, where each instance l satisfies Assumption 1:

$$J^l = \min_{s \in \mathbb{S}^l} J^l(s), \quad (\mathbb{S}^l, J^l) \in \mathbb{O}, \quad l \in L. \quad (10)$$

Furthermore, we assume that it is possible to randomly sample indices l from L , e.g. via the uniform distribution.

Example Setting: Here, \mathbb{O} is a set containing pairs of objective functions and decision spaces. To provide an example consistent with the sections to follow, consider the following nonlinear dynamical system with state x and input u :

$$x_{k+1} = f(x_k, u_k), \quad x \in \mathcal{X}, \quad u \in \mathcal{U} \quad (11)$$

Provided a cost function over states and inputs, state constraints, and torque bounds, we can construct the following finite-time optimal controller with horizon H for the aforementioned system. Here, all state constraints are projected to input constraints through prediction over the model (11):

$$\begin{aligned} & \underset{\mathbf{u}=(u^0, u^1, \dots, u^{H-1}) \in \mathcal{U}^H}{\operatorname{argmin}} && J(\mathbf{u}, x_k), && \text{(FTOCP)} \\ & \text{subject to} && && \mathbf{u} \in \mathbb{U}(x_k) \subseteq \mathcal{U}^H. \end{aligned}$$

The aforementioned finite-time optimal controller (FTOCP) collapses to the form in (10) if we consider an optimality set \mathbb{O} indexed by states $x \in \mathcal{X}$ — a specific form of indexing more generally referred to via $l \in L$ in (10).

Key Insight: The critical insight for this section then is as follows. If we were to randomly sample via a distribution π over \mathbb{O} , optimization problems of the form in (10) and calculate the optimality gap G^l of percentile solutions for

that problem, the corresponding gap is a sample of some real-valued random variable. By taking multiple independent samples of this random variable, we can provide, using recent results on risk-measure estimation, a probabilistic upper bound on this random variable, *i.e.* a probabilistic upper bound on achievable optimality gaps. Phrased formally, an existing theorem in the literature we will utilize is as follows.

Theorem 3. *Let X be a real-valued random variable with (perhaps) unknown distribution π , let $\{x_i\}_{i=1}^N$ be a set of N independent samples of X , and let x_N^* be the maximum such sample. $\forall \epsilon \in [0, 1]$, the following statement holds:*

$$\mathbb{P}_\pi^N [\mathbb{P}_\pi[x \leq x_N^*] \geq 1 - \epsilon] \geq 1 - (1 - \epsilon)^N.$$

Proof: This follows from a scenario argument, see [28], and is a re-phrasing of Theorem 4 in [25]. ■

To employ Theorem 3, we require the following definition.

Definition 2. Let $\mathcal{G}(N_p)$ be a real-valued random variable with distribution π_G and samples \mathbf{g} defined as follows: (1) Uniformly sample an index $l \in L$, (2) Take a percentile approach to solve (10) corresponding to this sampled index l , producing the solution $s_{N_p}^{l*}$, (3) Calculate and report as a sample \mathbf{g} , the optimality gap $G^l(s_{N_p}^{l*})$.

Then, we can upper bound the optimality gaps of percentile solutions to all optimization problems formed in the set \mathbb{O} , to some minimum probability. The formal statement will follow.

Theorem 4. *Let $\mathcal{G}(N_p)$ be as per Definition 2, and let $\{\mathbf{g}_i\}_{i=1}^R$ be a set of R independent samples of $\mathcal{G}(N_p)$ with \mathbf{g}_R^* the maximum such sample. Then $\forall \epsilon \in [0, 1]$, percentile solutions to optimization problems of the form in (10) will exhibit optimality gaps less than \mathbf{g}_R^* with minimum probability $1 - \epsilon$ and confidence $1 - (1 - \epsilon)^R$, *i.e.**

$$\mathbb{P}_{\pi_G}^R [\mathbb{P}_{\pi_G}[\mathbf{g} \leq \mathbf{g}_R^*] \geq 1 - \epsilon] \geq 1 - (1 - \epsilon)^R.$$

Proof: Stems via Theorem 3. ■

In short, if we wish to remove the secondary sampling requirement for the determination of optimality gaps, we need to be able to calculate the optimality gap for at least R independently chosen optimization problems of the form (10). Doing so permits us to make a general statement on the maximum achievable optimality gaps, to some minimum probability. Now, we will provide a few examples.

IV. EXAMPLES

A. Validating Theorem 2 — Traveling Salesman

The Traveling Salesman Problem (TSP) is a classic non-convex path planning problem referencing the identification of the path of the shortest distance traversing each node in a set once. Mathematically, consider a set of waypoints $W = \{w_1, w_2, \dots, w_{|W|}\}$ $w_i \in \mathbb{R}^2$ and the set of all paths over these waypoints $P = \{(i_1, i_2, \dots, i_{|W|}) \mid i_j \in \{1, 2, \dots, |W|\}, i_j \neq i_k, \forall j \neq k\}$. Then formally, the Traveling Salesman Problem is to

$$\min_{p \in P} \sum_{i=0}^{|W|-1} \|p_i - p_{i+1}\| + \|p_0 - p_{|W|}\|.$$

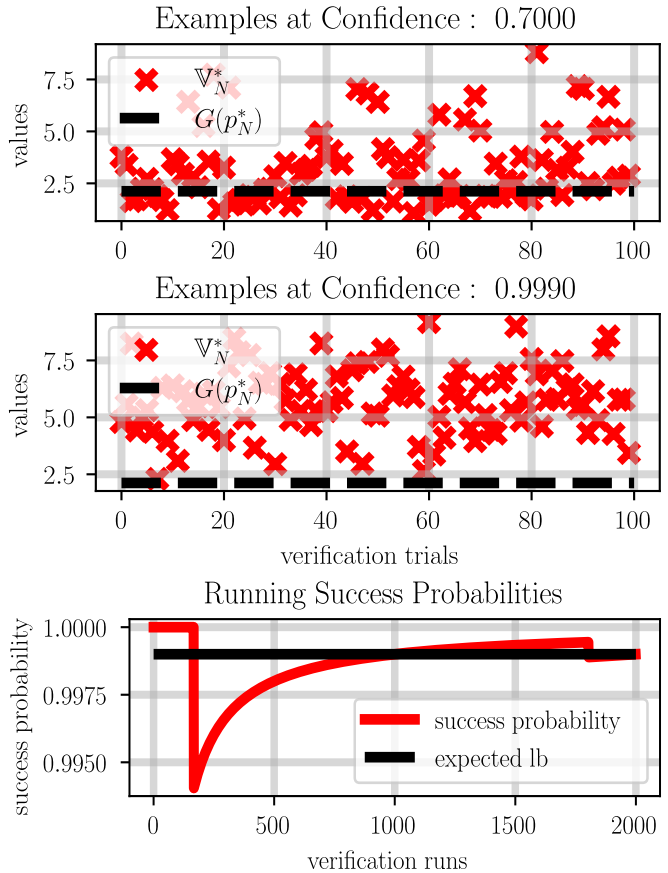


Fig. 2. Validation Data for Section IV-A corresponding to Theorem 2. (Top) 100 reported upper bounds $\mathbb{V}_{N_v}^*$ using Theorem 2 with desired confidence for (9) equal to 0.7. (Middle) 100 reported upper bounds $\mathbb{V}_{N_v}^*$ with confidence 0.999. (Bottom) Running fraction over 2000 trials of reported upper bounds $\mathbb{V}_{N_v}^*$ exceeding the true optimality gap $G(p_{N_p}^*)$ at confidence level 0.999. Notice how the fraction of upper bounds exceeding the optimality gap increases as we increase confidence (top to middle), and the running fraction of upper bounds exceeding the optimality gap converges to our desired confidence (bottom), corroborating Theorem 2.

For a graph with 10 nodes and 3628800 possible paths, evaluating $N_p = 5000$ paths and picking the best one identifies a path $p_{N_p}^*$ in the 99.9%-ile with at least 99% confidence according to Theorem 1. Using a subset D of the corresponding information set I_1 for the determination of such a percentile solution (see Assumption 2 for the definition of I_1), we define a variance function \mathbb{V} as per equation (4). Finally, to validate the probabilistic results of Theorem 2, we can also calculate the true probability p of uniformly sampling paths that exhibit a variance higher than the optimality gap of our proposed solution $G(p_{N_p}^*)$ — this is the minimum probability assumed to exist via Assumption 3 and defined in equation (7). For this particular node set and percentile solution $p_{N_p}^*$ the probability $p = 0.1083$.

Figure Analysis: To validate the results of Theorem 2, we solved for the minimum number of samples N_v required to determine an upper bound $\mathbb{V}_{N_v}^*$ to $G(p_{N_p}^*)$ with minimum confidence 0.7 — top figure in Figure 2 requiring 11 samples— and minimum confidence 0.999 — middle figure in Figure 2 requiring 61 samples. Both of these minimum sample requirements were identified by setting $\epsilon = p$ in (9)




Fig. 3. Validation data for Section IV-B. We claim that by varying the amount of information used to generate the variance function \mathbb{V} in equation (4), we can change the baseline probability p of sampling a decision whose variance exceeds the optimality gap of a given solution $G(s_{\chi}^*)$ (such decisions are highlighted in orange). Notice that as the volume fractions χ occupied by the chosen information set D decreases, we see a corresponding increase in the baseline probability p . Section IV-B discusses why this inverse relationship holds.

Name	N_p	N_v	expected success probability (9)	true success probability	average runtime (ms)
R-2	300	300	0.95	≈ 1	5.18
R-10	300	300	0.95	≈ 1	5.17
Ack	300	300	0.95	≈ 1	5.28
Ble	300	300	0.95	≈ 1	5.17
Levi	300	300	0.95	≈ 1	5.31
Himm	300	300	0.95	≈ 1	5.30

TABLE I

DATA FOR SECTION IV-B FOR THE (R-2) RASTIGRIN 2-D, (R-10) RASTIGRIN 10-D, (ACK) ACKLEY, (BLE) BEALE, LEVI, AND (HIMM) HIMMELBLAU BENCHMARK OPTIMIZATION PROBLEMS.

and solving for the minimum integer N_v required to make the right-hand side greater than or equal to our desired confidence. For each confidence level, we performed 100 separate trials and reported the maximum variance $\mathbb{V}_{N_v}^*$ produced by each trial according to Theorem 2. As can be seen in the corresponding data, increasing the confidence increases the likelihood that the corresponding reported result $\mathbb{V}_{N_v}^*$ exceeds the true optimality gap $G(p_{N_p}^*)$ — the red x 'es in the middle figure are all above the black, dashed line, whereas a few dip below the same line in the top figure when we report solutions with lower confidence. Furthermore, by repeating the procedure once more at confidence 0.999, taking 2000 separate verification runs, and recording whether $\mathbb{V}_{N_v}^* \geq G(p_{N_p}^*)$ per run, we can get a sense of the true, running probability that $\mathbb{V}_{N_v}^* \geq G(p_{N_p}^*)$. As can be seen in the bottom figure, this probability converges to 0.999 — the lower bound expected by Theorem 2. Notably, though, this result implies that we were able to identify a path in the 99%-ile that was no more than 1.12 times the length of the optimal path, by only evaluating 5061 uniformly chosen path samples, less than 0.14% of the overall decision space.

B. Increasing Success Probabilities — Benchmark Functions

In a brief remark after defining the variance function in equation (4), we mentioned that by specific choice of a subset D , one could increase the baseline probability p of uniformly choosing samples that exhibit a higher variance than the optimality gap of the reported percentile solution. This

section provides evidence in support of that statement for a few benchmark optimization problems. The one referenced in Figure 3, the 2-d Rastigrin function [29], is as follows:

$$\min_{x \in [-5.12, 5.12]^2} 20 + \sum_{i=1}^2 (x_i^2 - 10 \cos(2\pi x_i)). \quad (12)$$

Both the decision space and objective function are bounded, permitting a percentile solution to (12). Following Theorem 1 and taking $N_p = 100$ samples for such an approach, we generate the information set I_1 and percentile solution $x_{N_p}^*$. Then, by choice of a subset D of I_1 for the definition of the variance function in (4), we claim we can vary the baseline probability p of sampling decisions that exhibit a higher variance than the true optimality gap. To show this, define χ to be the volume fraction the subset D in (4) occupies of I_1 , i.e. $\chi = 1$ implies $D = I_1$ and $\chi = 0.05$ implies that D contains 1/20-th as many elements as I_1 .

Figure Analysis: Figure 3 portrays the results of varying the volume fraction χ of decisions D utilized to generate the variance function \mathbb{V} per equation (4). The decisions highlighted in orange are those that exhibit a higher variance than the optimality gap of the reported solution. Notice that as χ decreases, p increases as indicated in the titles. This inverse relationship arises as if we consider two sets $D_1 \subset D_2$ utilized for generation of the variance functions $\mathbb{V}_1, \mathbb{V}_2$, respectively, then $\forall s \in \mathbb{S}, \mathbb{V}_1(s) \geq \mathbb{V}_2(s)$ by definition of \mathbb{V} as a minimization problem. In other words, decreasing the amount of information provided to the variance function increases the corresponding conservativeness of the resulting function, which increases the likelihood of sampling a, now more conservative, upper bound on the optimality gap. In practice, and in the examples provided in the prior section, using a dilation $\chi = 0.1$ proved most effective, though studying if there exists an optimal volume fraction remains an open problem and the subject of future work.

Table Description: By increasing the success probability p of sampling decisions whose variance exceeds the optimality gap, we can “blindly” use Theorem 2 to bound the optimality gap of percentile solutions to benchmark optimization problems. Table I shows our data in this vein. For each

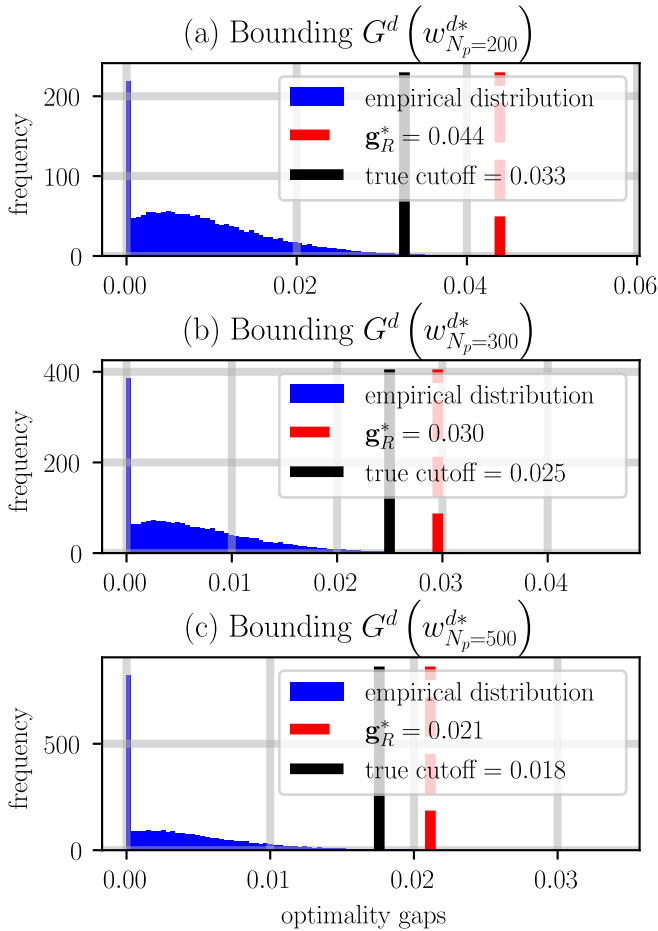


Fig. 4. Validation data for Section IV-C in support of Theorem 4. We claim that we can upper bound the optimality gap of successive applications of percentile methods to solve optimization problems of the form in (10). Shown above in red are the calculated upper bounds for the black lines corresponding to the 99% cutoff value of optimality gaps for percentile solutions to (NMPC-B). For the three separate percentile methods shown, we’re able to upper bound the true value every time, corroborating Theorem 4. This information is further discussed in Section IV-C.

benchmark optimization problem, we produced a percentile solution using $N_p = 300$ samples, constructed a variance function \mathbb{V} using 10% of the information set I_1 generated via the percentile method, and took $N_v = 300$ samples to identify the probabilistic maximum variance $\mathbb{V}_{N_v}^*$. Under the assumption that $p \geq 0.01$, Theorem 1 tells us that $\mathbb{V}_{N_v}^* \geq G(x_{N_p}^*)$ with 95% probability — column 4 in Table I. Indeed, over 5000 trials following the above procedure for each optimization problem, we were successfully able to identify a valid upper bound every time. Additionally, as the sampling method only requires the evaluation of sampled points, which is relatively quick, the procedure takes very little time to implement, as seen in the rightmost column.

C. Validating Theorem 4 — Nonlinear MPC

To validate the results of Theorem 4, we require a series of optimization problems of the form in (10). Keeping with the example provided in Section III-B, we aim to bound the optimality gap of percentile solutions for a Nonlinear MPC controller steering the Robotarium robots [30] — a collection of agents modelable via unicycle dynamics. Formally, let

$x \in \mathcal{X} = [-1.6, 1.6] \times [-1.2, 1.2] \times [0, 2\pi]$ be the system state, and let $u \in \mathcal{U} = [-0.2, 0.2] \times [-\pi, \pi]$ be the control input. Then, our model is as follows:

$$x_{k+1} = x_k + \underbrace{\begin{bmatrix} \cos(x_k[3]) & 0 \\ \sin(x_k[3]) & 0 \\ 0 & 1 \end{bmatrix} u_k}_{f(x_k, u_k)} (\Delta t = 0.033).$$

Furthermore, each agent has a Lyapunov controller U steering it to a provided waypoint $w \in \mathcal{W}$:

$$U : \mathcal{X} \times \mathcal{W} \triangleq [-1.6, 1.6] \times [-1.2, 1.2] \rightarrow \mathcal{U}.$$

We will use U to construct our MPC algorithm, which provides waypoints to steer the system around static and moving obstacles toward at least one goal. Formally, we overlay an 8×5 grid on \mathcal{W} and define the space of operating environments \mathcal{D} as those environments that: (1) have 8 static obstacles (SO) and 3 goals (g), (2) have controlled (x_a) and uncontrolled (x_o) agent starting positions outside static obstacles and goals, and (3) have at least one feasible path from the controlled agent’s starting location to one of the three goals. A vector $d = [x_a, x_o, \text{SO}, g] \in \mathcal{D}$ corresponds to one such environmental setup. To account for collisions, consider the following barrier function h , where $P = [\mathbf{I}_{2 \times 2} \ \mathbf{0}_{2 \times 1}]$ projects system states to the plane [7]:

$$h(x_a, x_o, d) = \begin{cases} -5 & \text{in SO cell,} \\ \|P(x_a - x_o)\| - 0.18 & \text{else.} \end{cases}$$

Then the nominal NMPC algorithm minimizes $S : \mathcal{W} \rightarrow \mathbb{R}$ — a function outputting the shortest path distance from a waypoint to the closest goal — while ensuring that the existing Lyapunov controller U renders the barrier h positive for the next 5 time steps, *i.e.*, $\forall j = 1, 2, \dots, 5$,

$$\begin{aligned} w_k^* &= \underset{w \in \mathcal{W}}{\operatorname{argmin}} && S(w), && \text{(NMPC-A)} \\ \text{subject to} &&& x_k^j = f(x_k^{j-1}, u^{j-1}), && \text{(a)} \\ &&& x_k^0 = x_k, && \text{(b)} \\ &&& h(x_{k,a}^j, x_o, d) \geq 0 && \text{(c)} \\ &&& u^{j-1} = U(x_k^{j-1}, w), && \text{(d)} \\ &&& 0.05 \leq \|w - x_k\| \leq 0.2. && \end{aligned}$$

To ease sampling requirements, we consider an augmented cost J that outputs 100 whenever a waypoint w fails to satisfy constraints (a)-(d) in (NMPC-A), yielding the following:

$$\begin{aligned} w^{d*} &= \underset{w \in \mathcal{W}}{\operatorname{argmin}} && J(w, d), && \text{(NMPC-B)} \\ \text{subject to} &&& 0.05 \leq \|w - x\| \leq 0.2. && \end{aligned}$$

Finally, we note that (NMPC-B) is equivalent to (10), if we consider as index set L , the environment set \mathcal{D} .

Figure Analysis: To validate the results of Theorem 4 we uniformly sample $R = 459$ different environments $d \in \mathcal{D}$ and calculate a percentile solution to (NMPC-B) with $N_p = 200$ samples, Figure 4-(a); $N_p = 300$ samples, Figure 4-(b); and $N_p = 500$ samples, Figure 4-(c). We calculate

the optimality gap for each solution by performing gradient descent on the best out of 2000 uniformly chosen samples and reporting the final value as the true optimal value. According to Theorem 4, in each case, we should produce an upper bound on optimality gaps \mathbf{g}_R^* that exceeds sample-able optimality gaps \mathbf{g} with at least 99% probability and 99% confidence. To validate this statement, we uniformly sampled 50000 more environments $d \in \mathcal{D}$ and followed the prior scheme for each percentile case to determine the distribution of sample-able optimality gaps. This data is reflected as the distributional data you see in each subfigure in Figure 4. As can be seen, in each case the reported upper bound \mathbf{g}_R^* exceeds the true 99% cutoff. Furthermore, as the number of samples taken for the percentile solution increases, the upper bound decreases. This is expected as we are providing a solution in a higher percentile each time. To emphasize the utility of this result for controls, say we wished to implement a percentile procedure with $N_p = 300$ samples to provide “good” waypoints optimizing for (NMPC-B). Offline calculation of this optimality gap would provide confidence that in practice, we would, with 99% probability and with 99% confidence, be choosing waypoints within $0.03m$ of the optimal waypoint at every iteration. Notably, we would not have to solve the non-convex program and repetitively sample variances at each time step to make this statement.

V. CONCLUSION

We provide a method to bound the optimality gaps of percentile methods to a wide class of non-convex optimization problems and explore the utility of such approaches for control in an NMPC setting. We also verify our mathematical results on several benchmark optimization problems including the traveling salesman problem. In future work, we hope to decrease the conservatism of our bounds and formally prove the existence of a baseline probability for certain classes of optimization problems, *i.e.* remove the requirement on our fairness assumption preceding our theoretical statements.

REFERENCES

- [1] F. L. Lewis, D. Vrabie, and V. L. Syrmos, *Optimal control*. John Wiley & Sons, 2012.
- [2] R. E. Kalman *et al.*, “Contributions to the theory of optimal control,” *Bol. soc. mat. mexicana*, vol. 5, no. 2, pp. 102–119, 1960.
- [3] A. Locatelli and S. Sieniutycz, “Optimal control: An introduction,” *Appl. Mech. Rev.*, vol. 55, no. 3, pp. B48–B49, 2002.
- [4] E. F. Camacho and C. B. Alba, *Model predictive control*. Springer science & business media, 2013.
- [5] J. B. Rawlings, “Tutorial overview of model predictive control,” *IEEE control systems magazine*, vol. 20, no. 3, pp. 38–52, 2000.
- [6] C. E. Garcia, D. M. Prett, and M. Morari, “Model predictive control: Theory and practice—a survey,” *Automatica*, vol. 25, no. 3, pp. 335–348, 1989.
- [7] A. D. Ames, X. Xu, J. W. Grizzle, and P. Tabuada, “Control barrier function based quadratic programs for safety critical systems,” *IEEE Transactions on Automatic Control*, vol. 62, no. 8, pp. 3861–3876, 2016.
- [8] X. Xu, P. Tabuada, J. W. Grizzle, and A. D. Ames, “Robustness of control barrier functions for safety critical control,” *IFAC-PapersOnLine*, vol. 48, no. 27, pp. 54–61, 2015.
- [9] R. Grandia, A. J. Taylor, A. D. Ames, and M. Hutter, “Multi-layered safety for legged robots via control barrier functions and model predictive control,” in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 8352–8358.
- [10] P. Raja and S. Pugazhenth, “Optimal path planning of mobile robots: A review,” *International journal of physical sciences*, vol. 7, no. 9, pp. 1314–1320, 2012.
- [11] I. Noreen, A. Khan, and Z. Habib, “Optimal path planning using rrt* based approaches: a survey and future directions,” *International Journal of Advanced Computer Science and Applications*, vol. 7, no. 11, 2016.
- [12] B. Riviere, W. Hönig, Y. Yue, and S.-J. Chung, “Glas: Global-to-local safe autonomy synthesis for multi-robot motion planning with end-to-end learning,” *IEEE robotics and automation letters*, vol. 5, no. 3, pp. 4249–4256, 2020.
- [13] H. Ma, J. Chen, S. Eben, Z. Lin, Y. Guan, Y. Ren, and S. Zheng, “Model-based constrained reinforcement learning using generalized control barrier function,” in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 4552–4559.
- [14] A. Agrawal and K. Sreenath, “Discrete control barrier functions for safety-critical control of discrete systems with application to bipedal robot navigation,” in *Robotics: Science and Systems*, vol. 13. Cambridge, MA, USA, 2017.
- [15] J. Zeng, Z. Li, and K. Sreenath, “Enhancing feasibility and safety of nonlinear model predictive control with discrete-time control barrier functions,” in *2021 60th IEEE Conference on Decision and Control (CDC)*. IEEE, 2021, pp. 6137–6144.
- [16] J. Zeng, B. Zhang, and K. Sreenath, “Safety-critical model predictive control with discrete-time control barrier function,” in *2021 American Control Conference (ACC)*. IEEE, 2021, pp. 3882–3889.
- [17] G. Williams, A. Aldrich, and E. A. Theodorou, “Model predictive path integral control: From theory to parallel computation,” *Journal of Guidance, Control, and Dynamics*, vol. 40, no. 2, pp. 344–357, 2017.
- [18] M. S. Gandhi, B. Vlahov, J. Gibson, G. Williams, and E. A. Theodorou, “Robust model predictive path integral control: Analysis and performance guarantees,” *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 1423–1430, 2021.
- [19] J. Tamimi and P. Li, “Nonlinear model predictive control using multiple shooting combined with collocation on finite elements,” *IFAC Proceedings Volumes*, vol. 42, no. 11, pp. 703–708, 2009.
- [20] —, “A combined approach to nonlinear model predictive control of fast systems,” *Journal of Process Control*, vol. 20, no. 9, pp. 1092–1102, 2010.
- [21] O. Andersson, M. Wzorek, P. Rudol, and P. Doherty, “Model-predictive control with stochastic collision avoidance using bayesian policy optimization,” in *2016 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2016, pp. 4597–4604.
- [22] A. Jain and M. Morari, “Computing the racing line using bayesian optimization,” in *2020 59th IEEE Conference on Decision and Control (CDC)*. IEEE, 2020, pp. 6192–6197.
- [23] Y. Vaupel, N. C. Hamacher, A. Caspari, A. Mhamdi, I. G. Kevrekidis, and A. Mitsos, “Accelerating nonlinear model predictive control through machine learning,” *Journal of process control*, vol. 92, pp. 261–270, 2020.
- [24] D. Limon, J. Calliess, and J. M. Maciejowski, “Learning-based nonlinear model predictive control,” *IFAC-PapersOnLine*, vol. 50, no. 1, pp. 7769–7776, 2017.
- [25] P. Akella, A. Dixit, M. Ahmadi, J. W. Burdick, and A. D. Ames, “Sample-based bounds for coherent risk measures: Applications to policy synthesis and verification,” *arXiv preprint arXiv:2204.09833*, 2022.
- [26] P. Akella, W. Ubellacker, and A. D. Ames, “Probabilistic guarantees for nonlinear safety-critical optimal control,” *arXiv preprint arXiv:2303.06258*, 2023.
- [27] P. Akella, M. Ahmadi, and A. D. Ames, “A scenario approach to risk-aware safety-critical system verification,” *arXiv preprint arXiv:2203.02595*, 2022.
- [28] M. C. Campi and S. Garatti, “The exact feasibility of randomized solutions of uncertain convex programs,” *SIAM Journal on Optimization*, vol. 19, no. 3, pp. 1211–1230, 2008.
- [29] L. A. Rastrigin, “Systems of extremal control,” *Nauka*, 1974.
- [30] S. Wilson, P. Glotfelter, L. Wang, S. Mayya, G. Notomista, M. Mote, and M. Egerstedt, “The robotarium: Globally impactful opportunities, challenges, and lessons learned in remote-access, distributed control of multirobot systems,” *IEEE Control Systems Magazine*, vol. 40, no. 1, pp. 26–44, 2020.